# Risk-Neutral Probabilities Explained

Nicolas Gisiger

MAS Finance UZH ETHZ, CEMS MIM, M.A. HSG
E-Mail: nicolas.s.gisiger@alumni.ethz.ch

**Abstract**

All too often, the concept of risk-neutral probabilities in mathematical finance is poorly explained, and misleading statements are made. The aim of this paper is to provide an intuitive understanding of risk-neutral probabilities, and to explain in an easily accessible manner how they can be used for arbitrage-free asset pricing. The paper is meant as a stepping-stone to further reading for the beginning graduate student in finance.

---

*Version*: 10th of October 2010; First working paper version: 22nd of April 2009

*Acknowledgment*: I am grateful to Jakob Gisiger and Stephen McKeon for valuable feedback on earlier versions of this paper.



Electronic copy available at: http://ssrn.com/abstract=1395390

„Gentlemen, you are now about to embark on a course of studies which will occupy you for two years. Together, they form a noble adventure. But I would like to remind you of an important point. Nothing that you will learn in the course of your studies will be of the slightest possible use to you in after life, save only this, that if you work hard and intelligently you should be able to detect when a man is talking rot, and that, in my view, is the main, if not the sole, purpose of education."

- John Alexander Smith, Professor of Moral Philosophy, Oxford University, 1914

## 1. Introduction

Mathematical finance is concerned with the pricing of redundant securities. A security is called redundant if its payoff can be matched by holding other securities that already exist in the market, so-called primary securities. The process of creating a portfolio in order to match the payoff of a redundant security is called replication. The redundant security and the replication portfolio will have exactly the same payoff. Two positions with the same payoff must also have an identical market value. If the value was different, then this difference could be locked-in as a risk-free profit by engaging in arbitrage. Arbitrage means to buy the cheaper position and to sell the more expensive position. A profit is made immediately and no risky exposure is left for the future since the payoffs of the two positions will precisely cancel each other. Mathematical finance tries to establish precise relationships between different securities by assuming that arbitrage activities do not exist. It is about relative pricing, i.e. the price of a security is always expressed in terms of other securities.

The prices of the primary securities themselves are assumed to be given by the markets and are not explained by mathematical finance. Financial economics, on the other hand is a wider field, it tries to explain the pricing of the primary securities as well, via concepts such as endowments, preference functions, etc. Mathematical finance is therefore a subfield of financial economics.

Mathematical finance makes in its efforts extensive use of the risk-neutral probability concept. This concept is so widely used, that an intuitive understanding of it should not be avoided. In order to create this intuition and allow for a deeper understanding, we have to start exploring the concept from a financial economics' perspective. Once we understand the economic interpretation of risk-neutral



Electronic copy available at: http://ssrn.com/abstract=1395390

probabilities, however, we will accept that the prices of the primary securities are determined by the market, and we will proceed with relative pricing.

Sections 2.1 – 2.6 introduce all the basic notions in a one-period model. The relationship between arbitrage-freeness and the uniqueness of state prices is explained. [Note: all notions will be introduced step-by-step throughout the paper.] Risk-neutral probabilities are defined in terms of state prices, and interest rates are introduced. Section 2.7 moves to a multi-period model, which allows section 2.8 to show that every normalised expected asset price path is a martingale, and the implications thereof. Sections 2.9 – 2.10 introduce geometric Brownian motion as a continuous-time price process example, and outline how an arbitrage-free pricing formula can be obtained by moving from a variance estimate to the risk-neutral probability distribution, and from there to a state price distribution. The main result is that the drift component of the original geometric Brownian motion is not part of the final pricing equation, but substituted with the risk free rate; this is of significant help when trying to calculate the arbitrage-free price of a replicable asset.

## 2. Risk-neutral probabilities explained

### 2.1 Basic framework

A very simple framework is sufficient to understand the concept of risk-neutral probabilities. Imagine an economy which is in a known state at time 0, and which can move to a number of possible, mutually exclusive states at time 1. There is only time 0 and time 1. For example, imagine an economy which is now, at time 0, in state #3 (where the number 3 indicates the level of economic activity) and which can move to state #1, 2, 3, 4 or 5 at time 1. State #5 represents the highest amount of economic activity, and #1 the lowest. The economy can only be in one of these five states at time 1, but it is unknown at time 0, which state it will turn out to be. Imagine that there exists a homogenous view on the probability for these states to happen. We therefore have five probabilities, one for each state.[1] The probabilities sum up to 1, because we know that at time 1 the economy must be in one of these five states. This is our initial framework. It is easy to generalise the framework to [n] possible states, but this would not add any clarity to our explanation. Five states are easier to visualise. We first assume that interest rates are zero, allowing intertemporal transfer of wealth at no cost or benefit in dollar terms. This assumption will later be relaxed. Figure 1 illustrates our basic framework.

[1] These are not risk-neutral probabilities, but real probabilities.

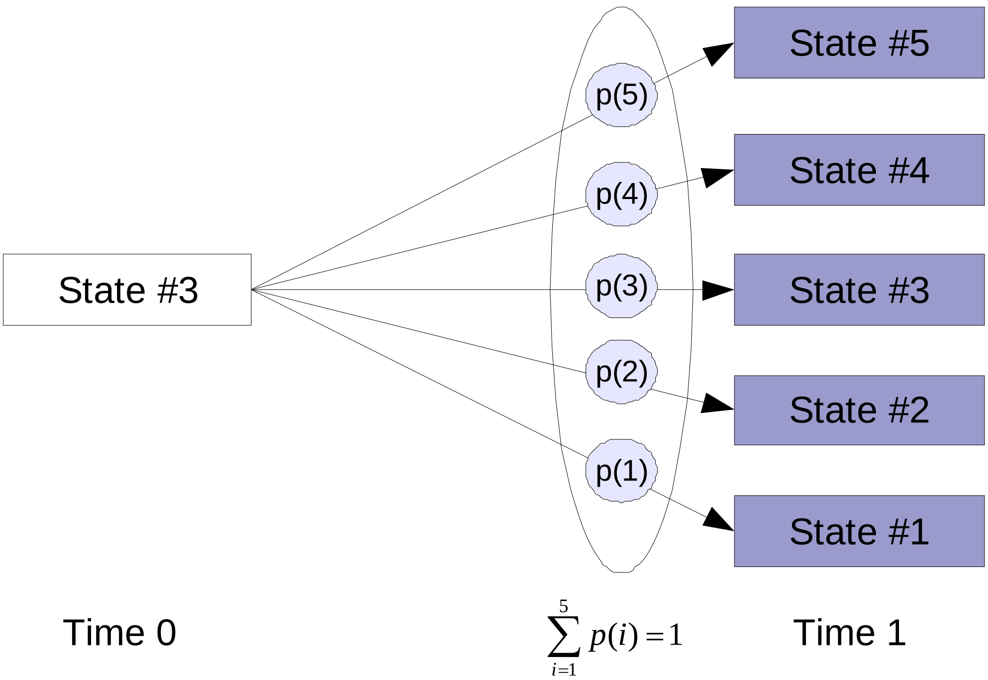


*Figure 1 – State space*

**2.2 Arrow securities**

Let us define five types of securities in this framework. For each state, we have one security with a contingent payoff of $1 if that particular state is reached, and $0 payoff otherwise. This is a so-called Arrow security. Figure 2 illustrates, as an example, the payoff of the Arrow security contingent on reaching state #4.

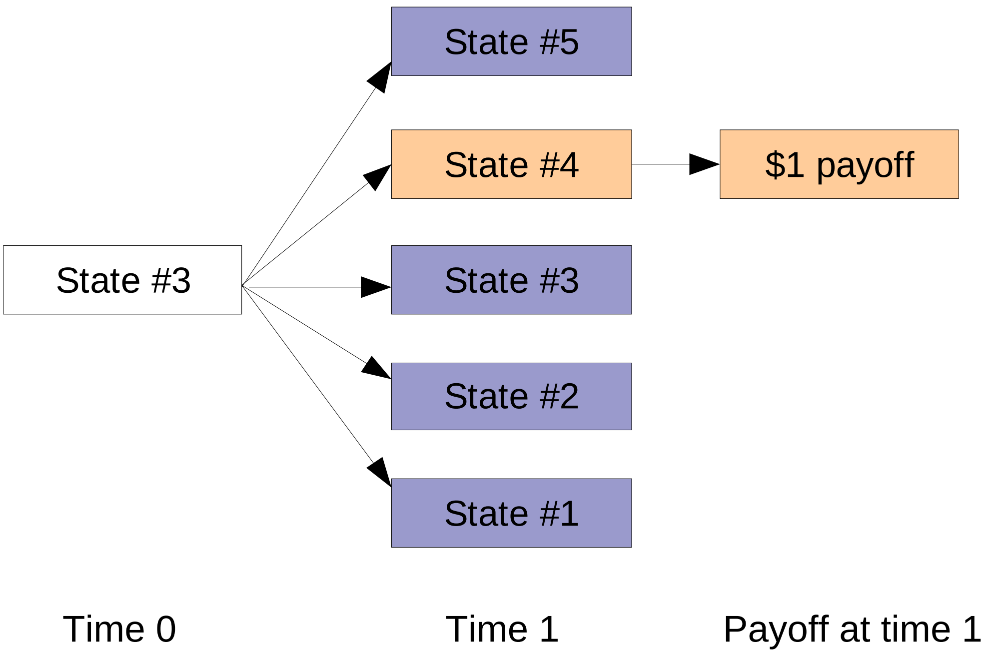


*Figure 2 – An Arrow security*

Therefore, if we hold all five Arrow securities at time 1, we are sure to receive a payoff of $1, since exactly one of the securities will have a payoff contingent on the reached state, and all the others will expire worthless. The value of a portfolio of all five securities is $1 at time 1. We presume the existence of a bank account with overdraft possibility. Since we assume interest rate for that bank account to be 0% in equilibrium, the arbitrage-free price of the whole Arrow portfolio must also be $1 at time 0. If the price of the portfolio was more than $1, say $(1+d), then we could sell such a portfolio in the market, thereby receiving $(1+d), and keep a sure profit of $d, since we will have to pay out exactly $1 at time 1. If the price was lower than $1, say $(1-d), we could buy such a portfolio at time 0,[2] thereby again locking in a sure profit of $d at time 1. By assuming that the equilibrium interest rate is 0%, these arbitrage activities will always drive the price of the complete set of Arrow securities towards $1.

What can we say about the price of each individual Arrow security? Each price will be determined by the supply and demand in the market. Relevant determinants of the supply and demand are: the *preferences* of the market participants with respect to holding money in one state versus another at time 1, the *preferences* with respect to holding money at time 0 versus time 1, and the estimated *probability* of a state actually occurring at time 1. If the participants deem $1 to be more valuable in state #1 than in state #5,[3] and both states are equally likely to occur, then the price of security #1 will be higher than the price of security #5. If the two prices were the same, then the participants would try to sell some amount of security #5 and buy security #1, because gaining some possible amount in state #1 would be worth more than giving up the same possible amount in state #5. This will push the market to price these two securities differently until the market is in equilibrium, i.e. quantity demanded of each security is equal to its quantity offered. If there was no preponderance of the value of dollars in one state over the other, but the probabilities of the two states were different, the same process would occur. If we do not care whether we have a possible dollar more in state #5 by giving up one possible dollar in state #4, and vice versa, but the probability of state #4 occurring was twice as high as state #5, then we would be willing to sell some amount of security #5 to get some more of security #4. It is the same idea as doubling the chances of winning in the lottery without incurring any additional cost. The perceived probability of each state will therefore positively impact the price of its linked Arrow security a time 0. Finally, by assuming that interest rates are zero, we are effectively implying that market participants have no preference of having $1 at time 0 over time 1, i.e. they are assumed to be indifferent with respect to time.

[2] With funding from the bank account
[3] It would not be surprising to value $1 more in times of a recession than in times of rapid economic expansion.

All we want to keep in mind at this point is that (1) the arbitrage-free price of the complete collection of Arrow securities at time 0 is $1, that (2) an Arrow security can command a premium over another security in line with the preferences of the market participants, and that (3) the more likely an Arrow security is to have a positive payoff, the higher its price at time 0. However, we do not have to calculate the individual prices but can simply observe them in the market, and specify a pricing vector **a** with the observed market prices. The price of security #1 at time 0 is a(1), the price of security #2 is a(2), etc.

We now move on to see if we can say more about a market where redundant securities are traded. A redundant security is a claim that has a payoff at time 1 which can be replicated by holding a linear combination of different Arrow securities. We need to distinguish *complete* markets from *incomplete* markets, since we have stronger results for complete markets.

### 2.3 Complete market

A complete market is a market where all Arrow securities can be traded. It does not matter whether there are any “real” Arrow securities or whether they can be constructed (i.e. replicated via a linear combination of other traded securities)[4]. As soon as we can construct every possible Arrow security, the market is complete. Remember that we know the price of each Arrow security, because we have observed all of them in the market.

We can now define a particular redundant security. A security is always defined by its, possibly stochastic, payoff. By *stochastic*, we mean that the payoff might depend on the state that the economy reaches at time 1. A risk-free zero coupon bond, for example, has a sure payoff of its notional [$n] at time 1, regardless of the economy’s state. In order to replicate the bond, we need to buy [n]-times the whole set of Arrow securities, as illustrated in figure 3. Since we already know that one whole set of Arrow securities costs $1 at time 0, the replication must cost $n. This is the only arbitrage free price of the bond. If the bond costs less than $n, say $(n-d), we could lock in a profit of $d by buying the bond and shorting [n]-times the whole set of Arrow securities. If the bond cost more than $n, say $(n+d), we could lock in a profit of $d by selling the bond and buying [n]-times the whole set of Arrow securities.[5]

[4] Such replication might involve long, as well as short positions, i.e. besides buying, we might have to sell some borrowed securities.

[5] In order to sell something we do not already own, we have to borrow it first. We enter into a repo transaction (repurchase agreement) with another party, i.e. we borrow the bond at time 0 and hand it back at time 1, including its payoff. This allows us to sell the borrowed bond at time 0, but we must buy it back in the market at time 1, including its payoff.

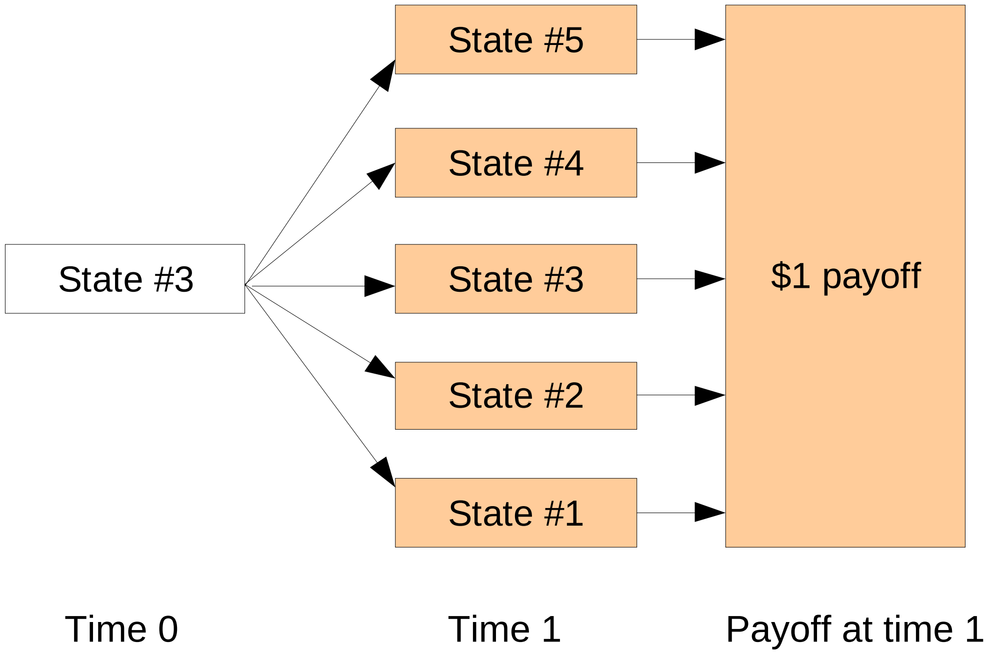


*Figure 3 – A risk-free zero coupon bond*

As a second example, a risky zero coupon bond has a stochastic payoff. Let us say, for example, that the risky zero coupon bond pays $1 in each state, except in state #1 where it pays nothing due to default. This is illustrated in figure 4.

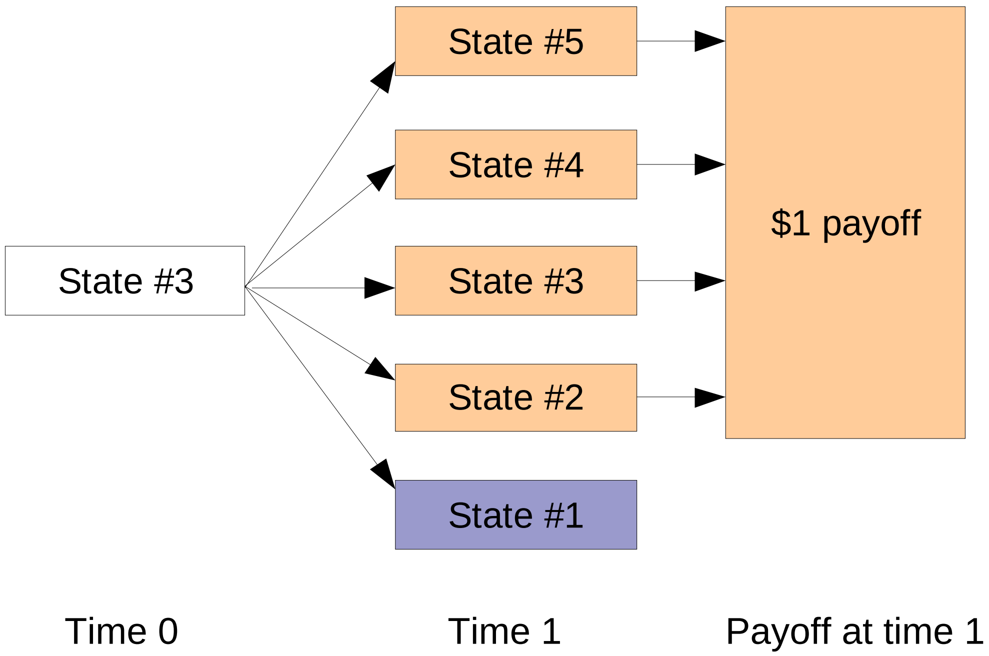


*Figure 4 – A risky zero coupon bond*

The only arbitrage free price of this bond is the sum of the prices of securities #2, 3, 4 and 5. If the bond was cheaper, say by $d, then we could buy it, sell Arrow securities #2, 3, 4 and 5, thereby locking in a sure profit of $d, because the long and short positions exactly cover each other at time 1, no matter in which state the economy will turn out to be. Similarly, if the bond was more expensive, then we would do the reverse operation; locking in the profit already at time 0 by matching the payoff of the long and short positions at time 1. Such arbitrage opportunities offer free profit; our assumption here is that such opportunities vanish because their pursuit pushes the market price towards the arbitrage-free level.

The idea seems very straightforward. **Any possible payoff can be replicated via a linear combination of Arrow securities; and this particular replication portfolio imposes a unique arbitrage-free price.** The arbitrage-free price is equal to the sum of all payoff-weighted Arrow security prices. If $\theta$ is the arbitrage-free price of a security at time 0, and $X(i)$ is the payoff of the security when state $i$ is reached, then we have:

$$\theta = \sum_{i=1}^{5} a(i) * X(i) \qquad (1)$$

Let us assume the following observed Arrow security prices: a(1) = 30 cents, a(2) = 25 cents, a(3) = 20 cents, a(4) = 15 cents, and a(5) = 10 cents, i.e. a security which pays $1 only if state #5 is reached at time 1, for example, costs 10 cents at time 0. Any payoff structure can be weighted with these prices in order to get the arbitrage-free price of the security at time 0. For example, the risky zero coupon bond that we presented before, has accordingly a unique arbitrage-free price of 70 cents, since it is composed of Arrow securities #2, 3, 4 and 5. We can replace the cents in the Arrow security prices with [%] if the payoff structure *X(i)* is already expressed in dollar terms. **These Arrow security prices are so-called risk-neutral probabilities; they are exactly the same thing.** It is surprising that this is not always pointed out very clearly. Note that we are currently still assuming interest rates to be zero.

Why should these prices be called *probabilities*? For a mathematician, a probability measure fulfils the criteria that the sum of all probabilities (of disjoint events) must equal 1, and that the probability of a particular event cannot be negative. Since the system of Arrow security prices outlined so far fulfils these conditions, it can technically be called a probability measure.[6]

[6] A *probability measure* is simply a mapping of outcomes to certain probabilities. In our example, each state of the economy has an assigned probability.

We have:

$$\sum_{i=1}^{5} a(i) = 1 \qquad (2)$$

and

$$a(i) \geq 0, \forall i \qquad (3)$$[7]

There are frequent statements about risk neutral probabilities which we can now easily identify as misconceptions. A risk-neutral probability is NOT the real probability of an event happening, but should rather be interpreted as a price. It is NOT INDEPENDENT of the real event probability, since the probability positively impacts the price of a state-contingent payoff. It does NOT assume that we are in a risk-free world, quite the opposite, the world is assumed to have an unpredictable future. It does NOT assume that market participants are risk-neutral, they can price different state-contingent payoffs in line with their risk preferences, be they risk averse, risk neutral or risk seeking. It is simply the only arbitrage free price in a complete market. **Risk-neutral probabilities (i.e. *Arrow security prices* or *state prices)*[8] simply enforce linear pricing consistency between all traded securities with regards to their payoff components.** The price for receiving $1 if state #3 occurs, for example, cannot be incorporated in a security's price at 25 cents in one case, and at 20 cents in another case. If two such securities would coexist, then this price difference could be singled-out by stripping both securities off all other stochastic payoffs;[9] and then trade these two stochastic payoffs (contingent on reaching state #3) against each other. Arbitrage activities are assumed to eliminate such pricing discrepancies.

### 2.4 Equivalent probability measures

The risk-neutral probability measure is *equivalent* to the real probability measure. A mathematician talks of the *equivalence* of two probability measures if both of them agree on the possible and the impossible outcomes. If one probability measure allocates a positive probability to a certain outcome (meaning that it is a possible outcome), then the other probability measure also allocates a positive probability to the same outcome. However, if the first measure allocates zero probability to an outcome (meaning that it is impossible), then the second measure also allocates zero probability. If these two conditions are satisfied, the two probability measures are called *equivalent*. This property of probability measures will be important when we try to derive a pricing formula later in this paper. We discuss it here already,

[7] $\forall i$ means *for all i*.
[8] We will use *Arrow security price* and *state price* interchangeably.
[9] With the help of Arrow securities.

since, if we think in terms of state prices, the equivalence of the two probability measures is easy to see.

If an outcome is impossible (according to the real probability measure) then the corresponding Arrow security cannot cost anything (risk-neutral probability). Similarly, if an outcome is possible (according to the real probability measure), then the price of the corresponding Arrow security must be more than zero (risk-neutral probability); otherwise we could get a probabilistic payoff for free. Arbitrage activities are assumed to push such a free security into a positive price range.

**2.5 Incomplete market**

In an incomplete market, not all Arrow securities can be constructed. A security which would be redundant in a complete market might not be so in an incomplete market. For example, we might not have enough traded securities in our market in order to trade or construct Arrow securities #1 and 2. All we can say is that a portfolio which contains securities #1 and 2 must have an arbitrage free price of [$1 - a(3) - a(4) - a(5)], because we know that the arbitrage-free price of all Arrow securities together sums up to $1. Now let us assume that a new security is introduced into the market which pays $1 if state #2, 3 or 4 is reached, but nothing in state #1 or 5. The payoff of this security is illustrated in figure 5.

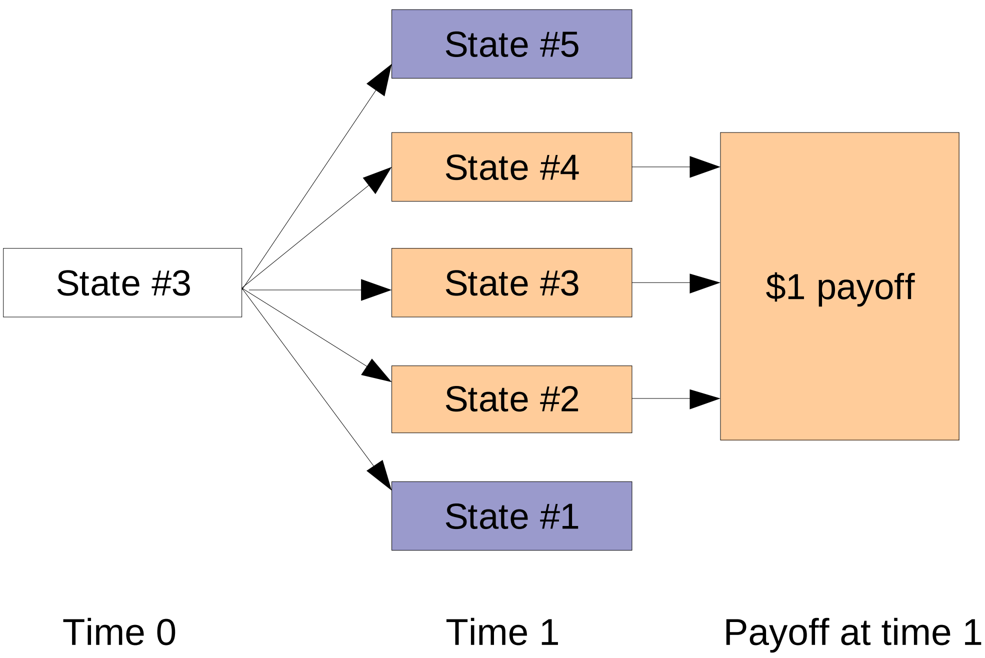


*Figure 5 – A security with stochastic payoff*

This security is not redundant, because we cannot strip-off all payoffs with a hedging portfolio. [10] The price of this security must be higher than [a(3) + a(4)], and we know that it must be less than [1 - a(5)]. Any price that is assumed within this range is arbitrage-free. This is the same as saying that the arbitrage-free price of Arrow security #2 is not unique; which in turn is the same as saying that the risk-neutral probability of state #2 is not uniquely determined. Several different risk-neutral probability measures can be accommodated by the market prices without leading to any arbitrage possibilities.

Once this particular security is trading in the market and a price is established, the market will be complete in our example. Any other new security would now be redundant and its price strictly determined by the other trading securities. Once we have a price for the security introduced before, we can determine a(2),[11] and therefore also a(1).[12]

We note that an incomplete market setting allows a range of arbitrage-free prices for certain securities. This is the same as saying that a range of risk-neutral probabilities exists with respect to the states for which no Arrow securities exist yet. Arbitrage-free pricing offers only limited guidance in such a setting. We will not look into incomplete markets any further.

**2.6 Relaxing the interest rate assumption**

So far we have assumed interest rates to be zero. We will now relax this assumption and observe what happens to state prices and risk-neutral probabilities, respectively.

Assuming an interest rate different from zero implies that the market allocates different utility to possessing $1 now versus later. $1 at time 0 can be *compounded* to [$1 * (1 + r)] at time 1. Similarly, $1 at time 1 can be *discounted* to [$1 / (1 + r)] at time 0. This interest rate is applied to our risk-free bank account.

We saw that $1 at time 1 is the same as the payoff of the complete set of Arrow securities. The price of the complete set of Arrow securities at time 0 must therefore be equal to [$1 / (1 + r)]; otherwise, one can make a risk-free profit by trading the complete set via funding through the bank account. We have:

$$\sum_{i=1}^{5} a(i) = \frac{1}{1+r} \qquad (4)$$

[10] The reverse position of the replication portfolio is called “hedging portfolio”. By *reverse*, we mean a long position is flipped to a short position, and vice versa.
[11] a(2) = [a(2) + a(3) + a(4)] – a(3) – a(4)
[12] a(1) = $1 – a(2) – a(3) – a(4) – a(5)

This discounting can be extended to all individual Arrow securities, i.e. the former price of each security is divided by [1 + r] to get the new price when introducing an interest rate different from zero. We assume that the risk preferences, state preferences and probability estimations remain the same; all that has changed in the market is time preference. Of course it could be envisaged that time preference is interrelated with other preferences, and that therefore state prices do not only change by the discount factor, but go through an additional distortion. All we can really be sure about is that the price of the whole set of Arrow securities is [$1 / (1 + r)] at time 0. In any case, we do not need to calculate the individual Arrow security prices; they are determined for us by the supply and demand in the market.

What happened to risk-neutral probabilities along the way? We know that the sum of the probabilities of all outcomes must be equal to 1, if we still want to maintain a mathematical probability concept. This is not the same as [1 / (1 + r)], if r is different from zero. We must therefore multiply each state price with [1 + r] in order to get the risk-neutral probability $q(i)$ for each state. As a sum, they again add up to 1. We have:

$$q(i) = a(i) * (1 + r), \forall i \tag{5}$$

and

$$\sum_{i=1}^{5} q(i) = 1 \tag{6}$$

One might argue that this starts to become an artificial construct; which is indeed true. A risk-neutral probability is simply something we define. However, as long as we remember that **risk-neutral probabilities are nothing else than compounded state prices**, we are fine. The reason for doing this construct is that there is value in keeping the condition alive that the sum of all individual probabilities adds up to 1. This means we can apply all the mathematical machinery that has been developed in relation with probability measures. We will see that this is useful at a later stage. For now, as long as we remember to reverse the artificial compounding with a subsequent equivalent discounting, we will always find our way back to state prices, the true determinants of each arbitrage-free price. In order to calculate the arbitrage-free price of a redundant security at time 0, we can simply multiply the payoff in each state with the corresponding risk-neutral probability, sum up all these products, and then discount this sum back to time 0.

$$\theta = \frac{1}{1 + r} \sum_{i=1}^{5} X(i) * q(i) \tag{7}$$

(7) is equivalent to equation (1). Any price diverging from this calculation would enable arbitrage, and market forces would push them to be in line again. In mathematical terms, we have done nothing else than discounting the payoff expectation under the risk-neutral measure Q. Formulae (8) and (7) are equivalent.

$$\theta = \frac{1}{1+r} E^{Q}\left[X\right] \quad (8)$$

We do now have an intuitive understanding of what risk-neutral probabilities are; they have a direct economic interpretation as compounded state prices. **Any factor that is relevant for determining the supply and demand of an Arrow security, and hence its price, has therefore a direct influence on the corresponding risk-neutral probability.** Any risk-premia implied by market prices are therefore incorporated in the risk-neutral probabilities. The frequently encountered statement that "*the risk-neutral probability is independent of the real probability*" is false. The real probability affects the state price, and is hence relevant for the risk-neutral probability. Of course, if we already know the state price, then there is no need to estimate the real probability for the sake of *calculating* the risk-neutral probability; as the estimate of the real probability is already incorporated in the state price.[13]

Thus far we have not explained why we rely on risk-neutral probabilities at all. Why should we calculate an arbitrage-free price via risk-neutral probabilities if we already know the state prices? The answer is that it might not be so straightforward to get to know the state prices; in reality, there is a continuous range of states., and therefore an infinite number of Arrow securities. Hence, we will later assume that state prices are unknown. **We will use the connection between risk-neutral probabilities and state prices the other way round i.e. we will obtain state prices from risk-neutral probabilities for the purpose of arbitrage-free pricing.** We will be able to use theorems from probability theory in order to obtain risk-neutral probabilities through a different route. This is the reason why we constructed an artificial probability measure. What follows in the next half of this paper is a brief overview of how we can calculate risk-neutral probabilities without using state prices. Before deriving an arbitrage-free pricing equation, however, we first have to gradually move towards a more realistic market setting, and introduce some further elements.

[13] A similar mistake is usually made when claiming that a forward FX rate is independent of the expected inflation in the two currencies, since the arbitrage-free forward FX rate is fully determined by the current FX rate and the two interest rates. Correctly, one would have to say that the current interest rates are affected by expected inflation. Hence, expected inflation is relevant for the arbitrage-free forward FX rate. However, we do not need to estimate expected inflation for the purpose of *calculating* the arbitrage-free forward FX rate, since this expectation is already reflected in the currencies' interest rates.

### 2.7 Trading strategy and dynamic completeness

So far, we have looked at a one period model only, with time 0 and 1. Conceptually, we can introduce a further element when moving to a multi period model. This further element is the *trading strategy*. In a one period model, all trading is a one-off decision at time 0, which cannot be changed before maturity of the security at time 1. Let us assume a simple multi-period model with times 0, 1 and 2, and two state branches at each time knot.

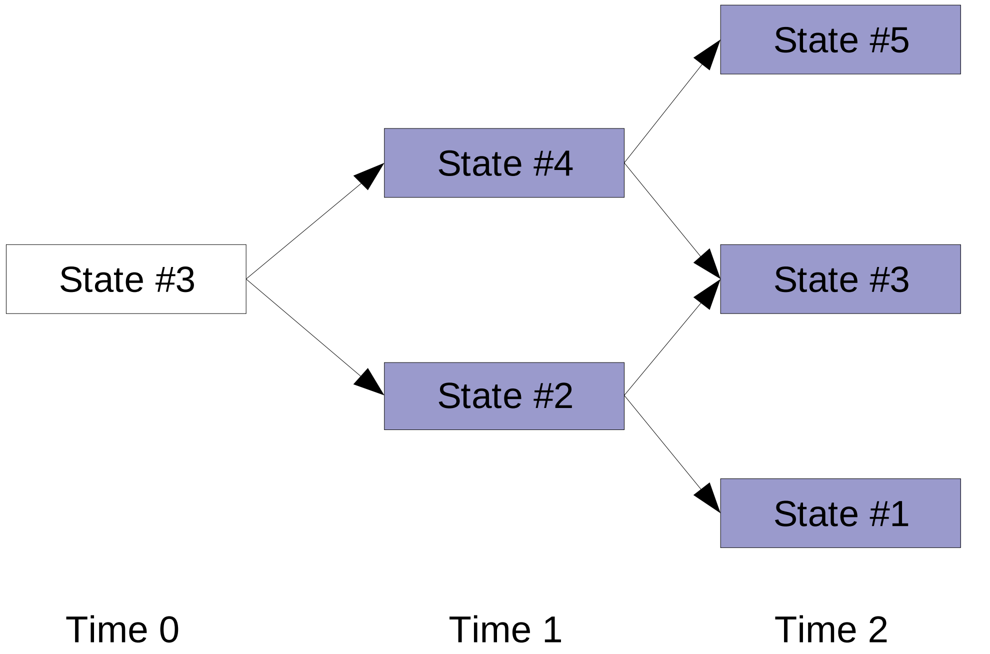


*Figure 6 – Multi-period state space*

In this model we are allowed to make trading decisions at time 0 and at time 1. A trading strategy can help in constructing a particular Arrow security.[14] If we can enter into a trade at time 0, adapt the trade at time 1 without incurring any additional cost or benefit , and be sure of the payoff at time 2 for each possible state, then the initial price of this future stochastic payoff is determined by the upfront cost of the initial trade. Any deviation from this price would lead to arbitrage activities. Basically, as soon as there is a replication strategy to perfectly match the payoff of a security at time 2, one can always hedge the position, thereby making sure that there is no exposure left at time 2. Price differences can be locked in at time 0 as a risk-free profit.

The possibility of intermediate trading does not make matters more complicated, quite the opposite; it makes it easier to construct Arrow securities. However, it is important to note that the trading strategy must be *self-financing*, i.e. there is no

[14] By *trading strategy* we are not referring to a strategy designed to make a profit, but to a replication strategy where we simply try to match a payoff.

money coming in or out before maturity. It is only about rebalancing the trading portfolio between its components at time 1 in a value-neutral way.[15] If the trading strategy was not self-financing, then we could not be sure of the initial cost to attain the final stochastic outcome,[16] and it would therefore not serve as a potential price enforcing arbitrage vehicle.

In the one-period model we required the availability of an Arrow security for each final outcome, which in turn requires at least an equal number of linearly independent securities in case we need to construct these Arrow securities first.[17] If we can adapt our trade through time, we can potentially construct such Arrow securities with fewer primary securities. A market, in which every Arrow security is available due to replication-enabling trading strategies, is called *dynamically complete*. A dynamically complete market determines the state prices and the risk-neutral measure, i.e. every redundant security has a unique arbitrage-free price.

For example, the Black-Scholes equation for calculating the arbitrage-free price of a European call option relies on a trading strategy whereby the replicating positions in the primary assets are continuously adjusted according to the *delta hedging rule*. Illustrations of such dynamic replication strategies can be looked-up in every book listed in the bibliography. Taleb (1997) should be consulted in order to become aware of the difficulties of dynamic replication.

For the concept of arbitrage-free pricing via risk-neutral probabilities, we actually do not rely on any knowledge about the exact replication strategy itself. We only rely on the fact that a replication strategy does exist. Of course, a replication strategy might not exist as often as we would like to assume in practice. All we want to keep in mind at this point is that replication might involve some active trading of the primary assets before maturity of the redundant asset, and not all payoffs can be reached by simply buying and holding a mix of primary securities from start until maturity. This trading should be *self-financing*, and the trading rule needs to be clear with the information available at the time of application. We cannot use the active replication strategy for arbitrage reasoning without these conditions.

### 2.8 Discounted asset prices as martingales

In a multi-period model with times 0, 1 and 2 we can write an asset price with a time subscript. $\theta_t$ is the price of the asset at time t. $\theta_t$ follows a stochastic process, adopting different values depending on the state of the economy at time t.

---

[15] Think of your own stock portfolio. You may decide to sell some amount of one stock to get into another one. If you are fully invested before and after your trade without taking any money out or injecting any more, then your trade is self-financing.

[16] We do not necessarily know at time 0 what the value of the security and the value of the replication portfolio are at time 2 (stochastic / uncertain). All we know is that they will be of equal value.

[17] Linear algebra: to solve for a number of unknowns, we need at least an equal number of equations.

Obviously, the price at time 2 is equal to the final payoff: $\theta_2 = X$ . It is important to distinguish between an actual price path and its expected path at time 0. In reality, a price will follow a certain path as uncertainty with respect to the payoff-relevant factors is resolved over time; this is the actual price observed over time represented by $\theta_t$ . However, we are here not interested in the actual path of the stochastic price process, because we simply cannot know what it will look like beforehand. The only thing we can do is taking expectations about the future at time 0.

To perform an expectation under a probability measure, the measure needs to be specified. We do not claim to be able to fully specify the real probability measure, i.e. the real probability of the price taking a certain path. However, let us assume that the risk-neutral probability measure is known. Therefore, we want to see how the price path behaves under the risk-neutral expectation operator, simply in order to discover its mathematical properties. We can only take expectations at time 0, because all the uncertainty is still ahead of us. This can be represented by:

$$E_0^Q[\theta_t], \forall t \geq 0 \tag{9}$$

This is simply a case of taking an expectation of an expectation, because the price itself is an expectation [as we saw in equation (8)]. The question is: what is our expectation at time 0 about our expectation at some point in the future? Let us start with the expected price at time 0:

$$E_0^Q[\theta_0] = E_0^Q\left[\frac{1}{(1+r)^2} E_0^Q[X]\right] \tag{10}$$

We simply replaced the price within the expectation operator with equation (8). Since we are now in a two period model, we are discounting for two periods. The discount factor is state-independent and can be taken out of the expectation operator. Finally, taking an expectation of the same expectation is the same as taking it only once. We therefore obtain:

$$E_0^Q[\theta_0] = \frac{1}{(1+r)^2} E_0^Q[X] \tag{11}$$

This is actually our arbitrage-free pricing equation (8). We could have written this equation directly. We have now shown what happens when taking the same expectation twice in a row, it is the same as taking it only once. We are now moving to the expected price at time 1, where we get:

$$E_0^Q[\theta_1] = E_0^Q\left[\frac{1}{1+r}E_1^Q[X]\right] \quad (12)$$

Since the price at time 1 is only one period before the payoff, we discount for one period only. Again, we can take the discount factor out of the expectation. Our expectation at time 0 of the final payoff, and our expectation at time 0 of our expectation at time 1 of the final payoff, must be the same. We do not know at time 0 what will happen to the state of the economy between time 0 and 1. This relationship is known as the *law of iterated expectations*.[18] We therefore obtain:

$$E_0^Q[\theta_1] = \frac{1}{1+r}E_0^Q[X] \quad (13)$$

Finally, we expect the price at time 2 to be:

$$E_0^Q[\theta_2] = E_0^Q\left[E_2^Q[X]\right] \quad (14)$$

At time 2, however, we will know what the outcome of the payoff is. The expectation of $X$ taken at time 2 must be the payoff itself, we can therefore write:

$$E_0^Q[\theta_2] = E_0^Q[X] \quad (15)$$

We can clearly recognise that the *expected* price grows at the risk-free rate $r$ with each period. This is not a surprising pattern. In fact, it is enforced by the arbitrage-free price of the risk-free zero-coupon bond (or bank account terms), which is traded for all maturities and grows at the risk-free rate $r$, and the way we defined risk-neutral probabilities. Relative to a state price, each individual risk-neutral probability is inflated until maturity of the payoff by the risk-free rate, simply because this growth rate is incorporated in its definition [see equation (5)].

Any *expectation* under the risk-neutral measure therefore grows at the risk-free rate. Therefore, the *expected price* of any asset under the risk-neutral measure grows at the risk free rate. That is, not the actual price will grow at the risk free rate, but the *expectation* of the price through time, where the expectation is calculated at a fixed time (time 0 in our example). This artificial drift was necessary to disguise

[18] Imagine that you have an expectation of the final score of a football match before the game starts. At the same time you can have an expectation about what your expectation of the final score will be after the first half of the game. When taking these expectations at the same time before the match, they coincide. It is only with passing time, as the game progresses and uncertainty is resolved, that these expectations can start to diverge, because then they are taken with a different set of information. The time subscript under the expectation operator effectively refers to the amount of information that we have available. When taking several expectations, we always have to work with the expectation that is based on the most restrictive set of information, which is obviously always the amount of information at an earlier point in time.

the state prices as a probability measure. We emphasize that we are not saying anything about how asset prices are actually developing over time as uncertainty is being resolved; all we know are the prices at time 0. We have now discovered an important mathematical property of our artificial probability measure.

We can get rid off this artificial expected drift only by going through the reverse operation. By doing so, every discounted expected future price[19] becomes equal to the price at time 0. We have:

$$\frac{1}{(1+r)^t} E_0^Q[\theta_t] = \theta_0, \forall t \geq 0 \tag{16}$$

We call this expression the *normalised expected price*. By going through this normalisation (i.e. discounting), we have achieved that the whole expression exhibits no more drift. By removing the drift, we have shown that the normalised expected asset price process is a *martingale*. A martingale is a stochastic process without expected drift. Since we can form a martingale out of every normalised expected asset price path with the help of risk-neutral probabilities, the risk-neutral probability measure is also called a *martingale measure*. The reason for transforming the expected price path into a martingale is to open the door for martingale theory. Any mathematical machinery and theories that have been developed for martingales can now be applied to arbitrage-free asset pricing.

The martingale property is very useful for our purpose of calculating arbitrage-free asset prices. We can simply turn equation (16) around to calculate the arbitrage-free price at time 0. This is the multi-period equivalent of equation (8):

$$\theta_0 = \frac{1}{(1+r)^t} E_0^Q[\theta_t], \forall t \geq 0 \tag{17}$$

In the next section, we will move to a framework where asset prices follow a continuous process. It can be shown mathematically that any *continuous square-integrable martingale* can be represented by a *Brownian motion* unfolding at a certain speed [Björk, 2004]. A Brownian motion is a stochastic process where each increment stems from a normal distribution. Hence, the transformation of asset prices into martingales with help of the risk-neutral measure effectively enables us to work with the normal distribution; but only under the restrictive assumption of finite variance. [20]

[19] From here on, whenever we talk of an expectation, we mean the expectation calculated with risk-neutral probabilities, i.e. taking a sum by weighting with compounded state prices.
[20] *Square-integrable* means that the variance of the process is finite. However, this assumption might not be accurate in reality, which would lead to a breakdown of our pricing framework.

### 2.9 Continuous-time model

We now move on to a setting where the price of a security is continuous in time and adopts values on a continuous range from 0 to $+\infty$. Security prices now follow continuous stochastic processes. Under the one-period or multi-period model we were assuming a discrete probability distribution for the price and final payoff. All that we are doing now is move to continuous probability distributions of the price and payoff. The probability of a price being at a certain level at a certain time can be characterised by a probability density function. The mathematics in a continuous setting are fairly advanced, but the key concepts developed so far remain all the same.

In reality, price processes, such as stock prices over time, are not continuous in any dimension, i.e. they are discontinuous in time and discontinuous in price. What do we mean? Transactions take place at a specific price, quantity and time. A transaction can therefore be represented by a point in a three dimensional space, characterised by price, quantity and time. The real price process is a sequence of transaction points. There is no connection between these points. The lines we usually draw through all the transaction points on a price chart are artificial. In reality, even if there was guaranteed continuous bid/offer quoting, a market maker might still not be able to execute at the quoted levels. It is therefore advisable to consider only transaction levels as price realisations, not quoted levels. We note that a continuous price process happens only in our model setting, not in the real world.

How can we specify the probability of an asset price to be at a certain level at a certain time, assuming that the process is now continuous? We choose the classic example of the geometric Brownian motion as a stock price model. The geometric Brownian motion is governed by the following *stochastic differential equation* (SDE):

$$dS_t = \mu S_t dt + \sigma S_t dW_t^P \tag{18}$$

$dW_t^P$ stands for the increment in a Brownian motion under a certain probability measure P. The increments of a Brownian motion follow a normal distribution, and this distribution is fully specified by its mean $E^P[dW] = 0$, and its variance $E^P[(dW)^2] = dt$. The parameter $\sigma$ scales the random shocks from the Brownian motion, $\mu$ specifies the deterministic drift through time. The solution to the SDE in (18) is given by:

$$S_t = S_0 \exp\left\{\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma W_t\right\} \tag{19}$$

This can be shown via an application of Ito's Lemma. We will here not get into the topic of solving SDEs, however, and simply present the result. Equation (19) specifies our probability distribution for $S_t$, as we know that $W_t$ follows a normal distribution: $W_t \sim N(0,t)$. Hence, the whole exponent $Z_t$ follows a normal distribution, since the other factors are only shifting and scaling the distribution in a non-random way:

$$Z_t = \left\{\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma W_t^P\right\} \tag{20}$$

$$Z_t \sim N\left(\left(\mu - \frac{\sigma^2}{2}\right)t, \sigma^2 t\right) \tag{21}$$

We will now try to give an arbitrage-free price to a redundant security defined by a payoff function $X$ that depends solely on $S_T$ (a certain stock price at time T). As a first approach, let us try to obtain the payoff expectation of the redundant security under the real probability measure. In order to do that, we have to sum up (over all possible scenarios) the payoff under each scenario with its corresponding probability density. As we know, normally distributed random variables take on values on the whole range from $-\infty$ to $+\infty$, our sum is therefore an integral. We get the following payoff expectation:

$$E_0^P\left[X(S_T)\right] = \int_{-\infty}^{+\infty} X\left(S_0 e^{Z_T}\right) p(\omega) d\omega \tag{22}$$

where $p(\omega)$ is the normal probability density function of $Z_T$.

Hence, we can calculate the real expectation of the security's payoff at maturity T, IF we have reliable estimates of $\mu$ and $\sigma$. We do know $S_0$ and T, and there are no further parameters. However, $\mu$ and $\sigma$ are very difficult to estimate objectively in reality, especially since we are interested in the *future* drift and volatility. How can we measure something that lies ahead of time? Moreover, what we are really interested in is the arbitrage-free price $\theta_0$, not the expected payoff at time T. Such a payoff expectation under the real probability measure has not been adjusted for any sort of risk premia such as those incorporated in the market prices of the primary securities. We have not shown that the above calculation leads to an arbitrage-free price. If we could somehow find a way of transforming the expectation from the real probability measure to the risk-neutral measure, then we could simply discount with the risk-free rate, and be sure that the obtained price is arbitrage-free. Such a probability measure transformation is based on the Girsanov Theorem.[21]

### 2.10 Obtaining the risk-neutral probability measure

We have learnt so far that risk-neutral probabilities are obtained by compounding Arrow security prices. We have in our continuous-time example not even made any attempt to replicate the redundant security's payoff, thereby finding out the required combination and amounts of Arrow securities and the required self-financing rebalancing strategy, which would then give us the arbitrage-free price. Is it possible to find out the arbitrage-free price in any other way? The answer is yes.

The Girsanov Theorem offers an alternative route. It shows how we can move a stochastic process driven by a Brownian motion from one probability measure to another, equivalent probability measure. Such a probability measure transformation is unique, and therefore, if we manage to obtain the dynamics under the risk-neutral measure, we can convert this to a state price distribution by simple risk-free discounting.

The move from one probability measure to an equivalent measure[22] is done via the Radon-Nikodym derivative $W(\omega)$, where $\omega$ represents a single scenario and $\Omega$ represents the set of all possible scenarios.

$$W(\omega)=\frac{q(\omega)}{p(\omega)}, \omega \subset \Omega \tag{23}$$

The density $q(\omega)$ is the risk-neutral probability density with respect to the state $\omega$, and $p(\omega)$ is the real probability density. The Radon-Nikodym derivative is the risk-neutral probability density with respect to the real probability density. We can now do the following transformation:

$$E_0^Q\left[X(S_T)\right]=\int_{-\infty}^{+\infty} X\left(S_0 e^{Z_T}\right) p(\omega) * \frac{q(\omega)}{p(\omega)} d\omega \tag{24}$$

[21] See for example Björk (2004)
[22] We have seen in section 2.4 that risk-neutral probabilities are equivalent to real probabilities under the absence of arbitrage possibilities.

Note that, as we change from one probability measure to another, the superscript above the expectation has changed from P to Q. This is the risk-neutral payoff expectation at time 0. Fortunately, we can easily move from the risk-neutral probability density to state price density $\varsigma(\omega)$ via simple discounting, just like in the discrete market setting. In a continuous setting, rather than thinking in terms of discrete Arrow security prices, it makes more sense to use state price density, since we would otherwise have to think in terms of an infinite number of Arrow securities distributed over a continuous range. **The existence of a unique risk-neutral probability measure implies the existence of a unique state price measure.**[23]

$$\varsigma(\omega) = e^{-rT} q(\omega) \tag{25}$$

The state price density is the market's fundamental pricing function with respect to all $\omega$. In a complete market, the arbitrage-free price of every asset must be consistent with the state price density, i.e. we have for every asset:

$$\theta_0 = \int_{\Omega} X(\omega) * \varsigma(\omega) d\omega \tag{26}$$

This is the continuous-range equivalent of equation (1). The state price density $\varsigma(\omega)$ is simply the equivalent of the Arrow security price in a continuous setting. The Radon-Nikodym derivative allows us, therefore, to obtain the arbitrage-free price of an asset. In our example, where the underlying asset follows a geometric Brownian motion, we have:

$$\theta_0 = e^{-rT} E_0^Q \left[ X(S_T) \right] = \int_{-\infty}^{+\infty} X\left(S_0 e^{Z_T}\right) p(\omega) * \frac{e^{-rT} q(\omega)}{p(\omega)} d\omega \tag{27}$$

So far, we know how to proceed conceptually, but how does $q(\omega)$ in this particular case look like? The Girsanov Theorem states that a Brownian motion under the real probability measure converts to a Brownian motion under an equivalent probability measure PLUS a drift component. We have:

$$dW_t^P = dW_t^Q + \varphi_t dt \tag{28}$$

We can plug this expression into the SDE in (18), and obtain the stock price dynamics under the new probability measure:

[23] This is based on the Fundamental Theorem of Asset Pricing.

$$dS_t = S_t(\mu + \sigma\varphi_t)dt + \sigma S_t dW_t^Q \tag{29}$$

However, we still know more than this. We know that the relative change in $S_t$ is composed of an expected drift of $r$ under the risk-neutral measure. The deterministic drift component arising from the measure change therefore precisely cancels the risk premium in $\mu$. The risk-neutral dynamics of the underlying asset are now fully specified with:

$$dS_t = S_t r dt + \sigma S_t dW_t^Q \tag{30}$$

We can directly move to these dynamics without ever knowing the specific risk premium. We know that the risk-premium must be cancelled, because we know the transformed drift already beforehand. After equation (27), we are left with the final arbitrage-free pricing equation for a redundant asset which is based on an underlying asset following a geometric Brownian motion:

$$\theta_0 = e^{-rT} \int_{-\infty}^{+\infty} X\left(S_0 e^{Z_T}\right) * q(\omega) d\omega \tag{31}$$

where $Z_T$ is normally distributed according to the normal density $q(\omega)$:

$$Z_T \sim N\left(\left(r - \frac{\sigma^2}{2}\right)T, \sigma^2 T\right) \tag{32}$$

$$q(\omega) = \frac{1}{\sigma\sqrt{2\pi T}} \exp\left\{-\frac{1}{2\sigma^2 T}\left[\omega - \left(r - \frac{\sigma^2}{2}\right)T\right]^2\right\} \tag{33}$$

**The nice aspect of this pricing approach is that we no longer need to estimate the parameter $\mu$. This is the whole merit of risk-neutral pricing.** The quality of our arbitrage-free pricing now depends on our ability to estimate $\sigma$.[24] How can we estimate this parameter?

The risk-neutral probability approach has helped us twofold in this respect. Firstly, we got rid of the drift component which leaves only one parameter to be estimated. Secondly, if we have access to the market prices of other securities which depend only on the same volatility parameter, thanks to our risk-neutral pricing formula,

[24] Do not forget that we have introduced a new parameter which also needs to be estimated, and this is the risk-free rate r. It is not necessarily clear in practice whether such a rate actually exists and which one it is. Nobody can rule out the possibility of a government default.

then we can calculate the implied volatility in these market prices, since it is the only unknown parameter. By calibrating our pricing formula with the implied volatility from other market prices we aim to get a market consistent, arbitrage-free price. Finally, whether this can be achieved depends largely on the validity of the assumptions inherent in our pricing framework.

# 3. Conclusion

### 3.1 Putting the pieces together

We have gone through many steps before finally arriving at an example of an arbitrage-free pricing equation. Here, we want to briefly summarise the chain of arguments again.

1. If a market is complete and arbitrage-free, then we have a unique system of state prices. The arbitrage-free price of any stochastic payoff is the sum (over all scenarios) of the scenario-payoff weighted with its state price. Based on the state prices, we can define a unique risk-neutral probability measure. For this purpose, we need to compound the state prices with the risk-free rate. This introduces a drift (at the risk-free rate) into any expectations under the risk-neutral measure taken at a fixed point in time. The arbitrage-free price of any stochastic payoff is therefore equal to the expected payoff under the risk-neutral measure, discounted at the risk-free rate.

This first chain of arguments has effectively linked arbitrage-free pricing with a probability concept. We can use this chain as a bridge, from arbitrage-free prices to equivalent probabilities, and vice versa. Most importantly, we can now use any tools and methods known in probability theory, and use this bridge to get back to arbitrage-free prices.

2. A redundant security's payoff is defined in terms of the price of an underlying security. If the uncertainty in the underlying security is driven by the innovation of a Brownian motion, then the Girsanov Theorem tells us how the price dynamics look like under the risk-neutral probability measure. This effectively allows us to obtain the risk-neutral probability distribution of the payoff that we are trying to price.

With the help of the second chain of arguments we have obtained the risk-neutral probability distribution of the redundant security's payoff. The first chain of arguments can now be used as a bridge to convert this probability distribution back to an arbitrage-free price.

### 3.2 Final thoughts

The main difficulty in intuitively grasping the risk-neutral pricing concept arises from the fact that the probabilistic methods and tools used were not primarily developed from a financial pricing perspective. This results in the use of two different languages, one from economics and the other from mathematics, which can easily be confusing.

We have shown in this paper that risk-neutral probabilities are, from the economist's standpoint, state prices compounded with the risk-free rate. One can always use this translation in order to make economic sense when using the tools from probability theory.

It should be mentioned that the risk-neutral probability concept is only useful for *arbitrage-free* pricing. An arbitrage-free price is not necessarily a *fair* price, or the *correct* price; it is only a *market consistent* price. We have two general conclusions: 1) if a market participant was buying (selling) a redundant asset above (below) its arbitrage-free price, then we can say that there would be a more efficient way for this market participant to express his view, namely via the replication strategy; 2) if a market participant was buying (selling) the underlying asset of a redundant security, where the redundant security trades below (above) its arbitrage-free price, then there would be a more efficient way for this market participant to express his view, again via the replication strategy.

Beyond these two statements, a unique arbitrage-free price only serves as a trading criterion if the market participant is ready to engage in arbitrage activities, trying to lock in price differentials via replication. Otherwise, one needs to take into account that the arbitrage-free price of a redundant asset is NOT independent of risk premia. The risk-neutral valuation approach implicitly uses the risk premia incorporated in the market price of the underlying asset. As we have seen, the price of the underlying asset is still part of the arbitrage-free pricing formula, therefore, the market risk premia still make their way into the equation. The market participant, however, might have different views on adequate risk premia than the market, and this needs to be taken into account when trading the redundant asset on a stand-alone basis.

Finally, we note that our pricing equation was not model-independent. What does this mean? The pricing equation in our example was still relying on the fact that a geometric Brownian motion was the correct model for the description of the underlying asset's dynamics. Plenty of simplifying assumptions were made. What the risk-neutral probability concept essentially helps us with is the specification of a model, by removing the need to specify the drift parameter. However, if the form of the model is not correct in the first place, then the specification thereof might not be of much use altogether. If the form of the model was different, then the so-called redundant asset might not be replicable and is therefore not redundant. All dynamic replication strategies are reliant on correct model assumptions. Making wrong assumptions about the form of the model leads us to the notion of *model risk*, which is probably by far the most delicate issue in quantitative finance. It would be much safer, for arbitrage-free pricing purposes, to rely on replication arguments which do not depend on a particular form for the probability distribution of the underlying asset price. It is very important to keep in mind that academics usually have a very strong tendency to strive for internal consistency when creating a model, even if this may come at the price of several simplifying assumptions. A practitioner who wants to apply such a model should never do so without checking external consistency, i.e. without ticking off and agreeing to each item in the assumptions list. In the best case, academics are close enough to practice in order to sufficiently satisfy external consistency in their models. In the worst case, a model which was simply meant to solve an idealized problem as an intellectual challenge is applied blindly under completely inappropriate circumstances.

* * * * * * *

The bibliography shows popular references for further introductory reading in mathematical finance. Especially Neftci (2000) is an excellent reference for an intuitive explanation of the theoretical underpinnings.

# Bibliography


BAXTER, MARTIN and RENNIE, ANDREW, *“Financial Calculus”*, Cambridge, UK: Cambridge University Press (1996)

BJÖRK, THOMAS, *„Arbitrage Theory in Continuous Time“*, Oxford: Oxford University Press (2004)

HULL, JOHN C., *„Options Futures and other Derivatives“*, 6th ed. Upper Saddle River, New Jersey: Prentice-Hall (2006)

NEFTCI, SALIH N., “An introduction to the Mathematics of Financial Derivatives”, 2nd ed. San Diego: Academic Press (2000)

SHREVE, STEVEN E., *„Stochastic Calculus for Finance I: The Binomial Asset Pricing Model”*, New York: Springer (2005)

SHREVE, STEVEN E., *„Stochastic Calculus for Finance II: Continuous-Time Models”*, 2nd ed. New York: Springer (2004)

TALEB, NASSIM N., *„Dynamic Hedging: Managing Vanilla and Exotic Options“*, New York: Wiley (1997)

WILMOTT, PAUL, *„Paul Wilmott on Quantitative Finance”*, 2nd ed. New York: Wiley (2006)